

A SUBDIVISION OPERATION AND USER INTERFACE FOR VOLUME SCULPTING

Field of the Invention:

This invention relates generally to methods of modeling three-dimensional objects and is also related to user interfaces that enable and facilitate the three-dimensional modeling methods.

Background of the Invention:

Sculpting has been shown to be a natural metaphor for the computer-modeling of three-dimensional objects. Hence, a variety of techniques that are analogous to sculpting have been proposed to generate three-dimensional computer models. Some of these techniques make use of small shapes or Constructive Solid Geometry (CSG) primitives that are agglomerated to form a model. Others techniques make use of volume elements, known as voxels, to build the models.

According to these techniques, primitives or voxels are either added to or taken away from the model. According to the carving technique, a large set of primitives or voxels is initially provided and primitives or voxels are gradually removed or "carved away" until a representative model of the object is formed. According to the marching tetrahedra technique, a small set of primitive is defined and repeatedly reproduced until all of the volume inside of an object is filled. Then, the primitives that lie outside of the volume of the object are removed, again leaving a model of the object.

One advantage of techniques analogous to sculpting is that precise numerical control of the vertices of the model, such as would be required for a CAD system, is not needed. A disadvantage of the techniques is that the preciseness of the model is limited by the size and shape of the primitives. Further, breaking down the primitives into smaller shapes, when possible, is problematic because many primitives must be broken down to avoid modeling inconsistencies and many voxels must be added to justify the presence of the added voxels.

Voxel-based techniques do not require the use of primitives with pre-determined shapes. According to voxel-based approaches, the user is presented with a grid of voxels and selects which of the voxels in the grid are inside of the object being modeled and which are outside of the object. Using binary values for the voxels, the voxels outside of the object can be turned off. Then, by forming polygons to contain the voxels that have not been turned off, isosurface approximations to the surfaces of the object being modeled can be extracted and output as a model.

1 When they are voxel-based, the carving and marching tetrahedra methods can use
2 a three-dimensional geometry, known as a kernel operator, to remove voxels from the
3 model. If voxels need to be added to the model, the kernel operator can perform a fill
4 operation that adds voxels to the model. However, because the size and shape of the
5 kernel operator typically remains fixed, the resolution of the model is limited.

6 In a related method called sawing, a cross-sectional area is extruded into three
7 dimensions and used as the kernel operator. Cross-sectional areas that are commonly
8 extruded according to the sawing method include circles, polygons, and Bézier curves.
9 However, the sawing operation can also be performed with any other closed path.
10 Nonetheless, the resolution of a model is still limited because the extruded kernel operator
11 is typically not subdivided.

12 In many three-dimensional modeling techniques, the model of the object can be
13 displayed in real time by a volume rendering technique. When it is voxel-based, the
14 model can be rendered as the set of voxels that have not been turned off or as the
15 isosurfaces of the polygons that contain those voxels. Once rendered, the voxels can be
16 "painted" to match the colors of the portions of the objects that they are representing.
17 However, the model is typically only displayed in a single view, making it difficult to
18 visualize the model in three dimensions.

19 As voxel-based modeling techniques have been developed, voxel-based interfaces
20 have also been developed and implemented in order to assist users in creating models.
21 Some interfaces include a three-dimensional cursor control device (mouse) and a set of
22 three-dimensional tools that can be used to remove or add voxels. However, most users
23 lack three-dimensional input devices and therefore require two-dimensional interfaces.

24 Although some two-dimensional interfaces and tools exist that provide sufficient
25 geometric intuition for sculpting three-dimensional objects, substantial training,
26 concentration, and practice is generally required of the users in order to use the tools. In
27 other words, currently available combinations of two-dimensional interfaces and tools for
28 sculpting objects in three dimensions are not intuitive and are therefore difficult to use.

29 **Summary of the Invention:**

30 A user interface for volume sculpting includes a processor, a two-dimensional
31 input device operably connected to the processor, and a display device operably
32 connected to the processor. The processor operates to provide on the display device a
33 first view of a sculpting object and a second view of the object. The first view of the
34 object provides a full, six-degree-of-freedom orientation control of the view.

1 A method of forming a model of a three-dimensional object includes generating a
2 three-dimensional set of points, grouping the points into a plurality of three-dimensional
3 cells, subdividing a cell in the plurality of cells into multiple subdivisions, locating
4 adjacent cells that contact the subdivisions, and subdividing the adjacent cells to eliminate
5 dangling points.

6 **Brief Description of the Drawings:**

7 The invention will be described by way of example, in the description of
8 exemplary embodiments, with particular reference to the accompanying drawings in
9 which:

10 FIG. 1 illustrates a perspective view of a representative model formed according
11 to a method of the present invention;

12 FIG. 2A, illustrate a first view a model as it is being formed. In this plan view, the
13 voxels that are included in the model are shown as thick points;

14 FIG. 2B illustrate a second view a model as it is being formed. This view is
15 oriented at a 90-degree offset to the first view and shows horizontal lines that indicate an
16 extrusion depth;

17 FIGS 3A-D illustrate a first cell that has been twice subdivided and an adjacent
18 cell that is subdivided according to one embodiment of the present invention to maintain
19 consistency of the model;

20 FIG. 4 illustrates the subdivision of a cubic cell into twelve tetrahedra;

21 FIG. 5 illustrates the subdivision of a cubic cell into six tetrahedra;

22 FIG. 6 illustrates a subdivision of the face of a cell that is adjacent to a cell that
23 has been subdivided into twelve tetrahedra;

24 FIG. 7 illustrates a subdivision of the edge of a cell that is adjacent to a cell that
25 has been subdivided into twelve tetrahedra;

26 FIG. 8 illustrates a flowchart of method of forming a three-dimensional model of
27 an object using a carving algorithm and/or a marching tetrahedra algorithm; and

28 FIG. 9 illustrates a system and user interface to execute methods according to the
29 present invention.

30 **Detailed Description of the Invention:**

31 FIG. 1 illustrates a perspective view of a representative three-dimensional model
32 of an object, a mug 10, as shown. The model can be viewed as a set of triangles that
33 make up the faces of primitives that contain voxels within the boundaries of the model.
34 The representative model in FIG. 1 can be formed by using modeling methods and

1 interfaces according to the present invention. The methods can include the marching
2 tetrahedra method and/or the carving method discussed above. Also, color, lighting,
3 and/or other effects can also be added to the model in order to give it some depth.

4 FIGS. 2A and 2B illustrate two views of the model, in this case a mug 10, as it is
5 being created. These views may be included in a user interface that allows a user to
6 implement certain methods of the present invention. FIG. 2A illustrates a first view that,
7 in this case, is a plan view of a cross-section of the model. In this view, the voxels
8 included in the model are shown as thick points. The first view can be moved with six
9 degrees of freedom and can therefore be moved or rotated fully relative to the model.
10 Thus, any perspective on the model that the user desires can be provided. FIG. 2A also
11 shows how the first view can include a coordinate system, as seen in the bottom, right-
12 hand corner of FIG. 2A.

13 The second view, illustrated in FIG. 2B, can be fixed in relation to the first view.
14 For example, as shown in FIG. 2B, the second view can be set to be orthogonal, or
15 substantially orthogonal, to the first view. When the first view is moved or redefined, the
16 second view can automatically adjust to remain orthogonal. Also, the distance from
17 which the model is viewed in the second view can be fixed.

18 Maintaining the orthogonal position of the second view relative to the first view
19 can be done by pre-loading the matrix stack of the modeled object with an offset matrix,
20 then drawing the same object with the same transformations as used in the first view.
21 Also, silhouette artifacts from the surface extraction method can be reduced or eliminated
22 by additional subdivision of the cells.

23 The horizontal lines 20, 30 illustrated in FIG. 2B indicate an extrusion depth and
24 can be set by a slider (not shown) in the user interface. The horizontal lines 20, 30 can be
25 used to define the depth range in which a carving, filling, adding, and/or painting
26 operation takes place.

27 Regardless of whether a marching tetrahedra, carving, or other modeling
28 algorithm is used, the first step of forming a three-dimensional model can include
29 generating a three-dimensional grid of voxels. These voxels can then be enclosed by, or
30 grouped within, one or more three-dimensional polygons, primitives, or cells. Cubic cells
31 will be discussed below, but cells of other geometries are also within the scope of the
32 present invention.

33 The initial number of voxels in the grid can be entered by a user or can be pre-set
34 in the modeling algorithm. Either way, the number of voxels in the grid sets the initial

any of the voxels in the adjacent, non-subdivided cell 80 and were therefore dangling voxels. In order to eliminate the dangling voxels, certain embodiments of the present invention create a new voxel at the centroid of the adjacent cell 80 and subdivide the adjacent cell 80 by joining the centroid voxel to the new voxels on the face of the adjacent cell 80 and to the corner voxels of the adjacent cell 80. This subdivision is shown in FIG. 3D and forms a set of pyramids, each having an apex at the centroid of the non-subdivided cell 80 and a square base with corners at four voxels.

In addition, FIG. 3D illustrates two more voxels that have been added at the midpoints of the edges of the adjacent cell 80. These two voxels are joined to the centroid voxel of the adjacent cell 80 and to the voxel in the center of the face of the adjacent cell shared with the eight even smaller cells 70. This joining forms additional pyramids within the adjacent cell 80.

The subdivision operation shown in FIGS. 3A-D allows for the elimination of dangling voxels in the cell 40 while only having to subdivide the adjacent cell 80 a single time. Previously available methods, on the other hand, would have required the adjacent cell 80 to be subdivided multiple times and would have required any cells adjacent to the adjacent cell 80 to also be subdivided to ensure that no dangling voxels were present in the adjacent cell 80.

According to the embodiment of the present invention illustrated in FIGS. 3A-D, the decomposition is such that only one new voxel needs to be added to the adjacent cell 80 in order to eliminate all dangling voxels. Hence, because only one subdivision operation is needed and because the total number of voxels is less than if the adjacent cell 80 had been twice subdivided, the process is simplified. The process is further simplified by the fact that no cells beyond the adjacent cell 80 have to be subdivided.

If the cell 40 had been subdivided more than twice, the adjacent cell 80 would still have required only one subdivision. In such cases, a centroid voxel would again have been added to the adjacent cell 80 and pyramids would again have been created. However, more pyramids, with smaller bases, would have been created.

In addition to the method discussed above, a cell can be subdivided by a variety of other methods. For example, a first alternate method is illustrated in FIG. 4 and forms a set of pyramids 60, each with five vertices, from a cubic cell (not shown). This method uses each of the faces of the cell as a base and the centroid of the cell as the apex of the pyramids 60. Then, each of the cell faces is divided into two right triangles, inducing in each pyramid 60 two tetrahedra. Thus, twelve tetrahedra can be formed in a cubic cell, as

1 shown in FIG. 4. These tetrahedra are non-overlapping, have triangular faces, and the
2 vertices created along their edges are consistent for all cubic cells that share that edge
3 (e.g., cubic cells adjacent to the cubic cell that was subdivided).

4 A second alternate method of subdividing a cubic cell is illustrated in FIG. 5 and
5 includes subdividing a cell (not shown) into six tetrahedra 90. This method subdivides
6 each face of the cell into two right triangles. Then, using a vertex on the other side of the
7 cube, this method creates six, non-overlapping tetrahedra.

8 A third alternate method of subdividing a cell calls for a cubic cell to be
9 subdivided into five tetrahedra. This method is analogous to the methods already
10 discussed and is therefore not illustrated or discussed further at this time.

11 Regardless of which of the illustrative methods of subdividing a cell discussed
12 above are used, each of the tetrahedra, pyramids, voxels, and cells can be associated with
13 a hierarchical level of subdivision based on how many subdivision operations they have
14 undergone. When a voxel and an adjacent cell are two or more levels of hierarchy apart,
15 a subdivision operation according to the present invention may need to be performed to
16 eliminate dangling voxels.

17 The subdivision method that forms twelve tetrahedra creates tetrahedra that have a
18 small aspect ratio and that are more likely to lead to small aspect ratios in the final
19 triangles of the model. Among the disadvantage of choosing a method that creates twelve
20 tetrahedra are that the method will almost surely create more triangles and that it will
21 require further interpolation of the volume properties of cells at their centroids. However,
22 this latter concern is minor, since linear interpolation can be used both in creating the
23 centroid and in creating the final triangle vertices.

24 The subdivision operations discussed above can be used in conjunction with the
25 previously mentioned marching tetrahedra and carving algorithms. In the marching
26 tetrahedra algorithm, the cells that are partially outside the volume being modeled can be
27 subdivided and the subdivisions that remain outside of the model can be selectively
28 removed. Because entire cells no longer need to be removed, the resolution of the model
29 is enhanced. Similarly, in the carving algorithm, the cells that lie partially outside of the
30 model can be subdivided and subdivided portions can be selectively removed.

31 The methods of the present invention not only give the user the opportunity to
32 remove portions of the model on an effectively arbitrary level of detail, but also give the
33 user greater flexibility in concentrating effort on the parts of the model that are of
34 particular interest at a given time. While the elimination of dangling voxels previously

1 required that many cells or primitives in the model be subdivided, according to certain
2 embodiments of the present invention, only the cells immediately adjacent to the dangling
3 voxels now require subdivision. Also, when used in conjunction with a user interface that
4 provides views such as, but not limited to, those illustrated in FIGS. 2A-B, a user can
5 view the model continuously as it is being formed and only the surfaces of the cells
6 affected by the subdivision operations need to be updated.

7 While the use of smaller cell subdivisions enhances the overall resolution of the
8 model, the possibility of locally subdividing cells allows for the presence of larger cells in
9 portions of the model where no small features are present. Hence, the total number of
10 voxels and cells in the model can be reduced or minimized.

11 According to certain embodiments of the present invention, both a voxel grid and
12 a hierarchical structure of cells is sometimes maintained. In such embodiments, carving
13 can be done on the voxel grid while subdivisions and surface extractions, such as the
14 marching tetrahedra method, can be performed by traversing the cell hierarchy. The same
15 is true of filling operations that put additional voxels and/or cells into the model and
16 painting operations that give the model color.

17 Both the voxels and cell data structures can track their level of subdivision within
18 the hierarchy. For example, a non-subdivided cell and all voxels within it can be placed
19 at a top level of hierarchy. A once-subdivided cell and all voxels within it would then be
20 one level below the non-subdivided cell, and so on. During surface extraction, when a
21 cell has a mixture of voxels that are at the hierarchical level of the cell level, such as the
22 voxels in the center of the cell, and voxels that are below the cell's level in the hierarchy,
23 such as the voxels on a face of the cell, the cell can be identified as being adjacent to a
24 subdivided cell. Once such an identification has been made, according to certain
25 embodiments of the present invention, a new subdivision operation can automatically be
26 triggered.

27 Any original subdivision scheme can be adapted to the present invention. For
28 example, a subdivision scheme that maintains consistent subdivisions between adjacent
29 cells can be used in conjunction with the method that forms six tetrahedra per cell. This
30 subdivision scheme can use a rotation scheme to maintain consistent ordering of the
31 vertices across cells. The twelve-tetrahedra-per-cell method can also be used, taking
32 advantage of the inherent symmetry of the subdivisions.

1 There are, however, two cases that require special handling of adjacent, non-
2 subdivided cells. The first case has to do with an adjacent face that is subdivided and the
3 second case has to do with an adjacent edge that is subdivided.

4 In the first of these cases, illustrated in FIG. 6, for each tetrahedron with a face in
5 a subdivided face, the voxels created by subdivision can be connected to the apex of the
6 tetrahedron and to the diagonally adjacent voxel within the same tetrahedron. Since each
7 tetrahedron has exactly one face coincident with a cell face, local control of changes to
8 volume can be maintained.

9 For adjacent edges that are subdivided, extra bookkeeping is necessary. Starting
10 with the original decomposition, a subdivided edge can split the original tetrahedron in
11 which it lies into two tetrahedra, as shown in FIG. 7.

12 If multiple edges are subdivided, considering them in sequence creates a set of
13 non-overlapping tetrahedra with voxels (and the centroid of the cell) as vertices. The
14 subdivision scheme can therefore carefully order the edges (in this case, from 0 to 7), as
15 shown in FIG. 7. In this example, two edges 420, 430, both in tetrahedra 400, 410 of two
16 adjacent cells (not shown), are subdivided and voxel 0 for the upper cell corresponds to
17 voxel 4 for the lower cell, etc. When subdividing, the upper cell examines the "2-1" edge
18 first and the lower cell examines the "5-6" edge first. The ordering creates in the shared
19 face the same tetrahedral (and model) edges. Then, the second voxel, introduced by
20 subdivision in the lower tetrahedra, uses the first new voxel (in the top tetrahedra) to
21 complete the subdivision operation and this algorithm again maintains local control.

22 The carving operations in the known volume sculpting method can be adapted to
23 the present invention. For example, to achieve a tapered three-dimensional operator, a
24 cross-section can be used as a wide cross-section and the edges can be tapered from the
25 middle of the extrusion range to the ends for an ellipsoidal operator. Alternately, the
26 edges can be tapered from the near to the far boundaries of the extrusion range for cone or
27 pyramid operators. Thus, certain embodiments of the present invention can approximate
28 a CSG-based modeler.

29 Certain embodiments of the present invention can help to infer an appropriate
30 function value at the voxels that are near a user-defined boundary. This can provide
31 smoother surfaces from operators that do not or cannot operate along a rectilinear voxel
32 grid. This also can reduce artifacts that are sometimes seen in the silhouettes of three-
33 dimensional models.

FIG. 8 illustrates a flowchart that defines the steps of modeling methods according to several embodiments of the present invention. Step 300 specifies that a three-dimensional set of points can be generated at the outset of modeling. The points can be chosen to be, among other items, voxels and the exact number of points or voxels generated can be pre-set or can be input by the user. Further, at a later time, more voxels can be added, should they become necessary.

Step 310 specifies that the voxels be grouped in a plurality of three-dimensional cells. Although, many modeling primitives with many different geometries may be used, the cubic voxels discussed above are well suited to certain embodiments of the present invention.

Step 315 calls for a user to input which cells to subdivide and step 320 calls for the cells chosen by the user to be subdivided into multiple subdivisions. As discussed above, a cell may be subdivided according to the methods illustrated in FIGS. 3A-D, and/or 4-7, wherein a cell may be divided into, for example, six, eight, or twelve subdivisions. Further, any other method that suits the needs of a user to obtain desired modeling resolution is within the scope of the present invention.

Step 330 specifies that cells adjacent to the subdivided cell be located. The adjacent cells typically share at least a portion of a face or edge with the subdivided cell and are typically at least one level of geometric hierarchy away from the subdivided cell.

Step 340 calls for the adjacent cells located in step 330 to be subdivided. According to certain embodiments, such as those illustrated in FIGS. 3A-D, the adjacent cells each only have one point, or voxel, added to them. However, the single-point-per-adjacent-cell addition can be enough to eliminate all dangling points in the model. Hence, this step can make it unnecessary to subdivide cells beyond the adjacent cells since all dangling points have already been removed. Further, according to certain embodiments, the subdivision of the adjacent cells can be performed automatically (i.e., without requesting input from a user).

Step 350 is a decision block that assesses whether the modeling technique being used is a carving algorithm or a marching tetrahedra algorithm. In the carving algorithm, all of the cells needed to model an object are generally included at the outset. Hence, all that remains is to subdivide cells and to remove subdivisions and/or cells. In the marching tetrahedra algorithm, on the other hand, the initial set of cell is generally formed into a suitable template, reproduced iteratively, and the reproduced templates are added to

1 the model until all of the volume within the object being modeled is filled. Then, portions
2 of the templates that lie outside of the volume are carved away.

3 Steps 360 through 380 detail how some carving algorithm embodiments are
4 implemented. According to step 360, which subdivisions to remove or fill is determined
5 and, according to step 365, at least one of the subdivisions is either removed or filled.
6 According to certain embodiments of the present invention, cells can also be removed if
7 none of the volume that they occupy is desired in the model. Step 370 then asks if the
8 model is complete pursuant to the performance of step 360. If the model is as
9 representative of the object being modeled as the user desires, then the algorithm stops.
10 On the other hand, if subdivision of additional cells within the model and/or the removal
11 or filling of additional subdivisions would make the model more representative of the
12 object, then there is either a return to step 360 or step 320. Whether the next step is step
13 320 or step 360 is determined at decision block 380, which determines further subdivision
14 of cells is necessary.

15 In the marching tetrahedra algorithm, represented by steps 390 through 420, step
16 390 specifies that a determination be made, within the algorithm, concerning where to
17 add subdivisions to the model. Then, step 395 calls for the addition of any necessary
18 subdivisions to the model. Step 420 then asks whether the model is complete. If the
19 model is complete at that point, the modeling process ends. However, if the model is not
20 complete, decision block 420 loops back to step 320, which subdivides cells and
21 eventually leads to further subdivisions and removals of cells and/or subdivisions.

22 In addition to the steps illustrated in FIG. 8, certain embodiments of the present
23 invention include the step of painting the points or voxels of the model. According to
24 these embodiments, color values can be assigned to each point or voxel, thus adding to
25 the clarity of the model or making the model more aesthetically pleasing. For example,
26 points may be shaded to represent depth or colors from a photograph of the object being
27 modeled may be superimposed on the model.

28 The embodiments of the present invention that have been discussed thus far have
29 been related to methods of generating voxel-based models of objects. The embodiments
30 discussed below, on the other hand, relate to simplified user interfaces for generating the
31 models. These interfaces can operate in three dimensions and can allow users to view a
32 three-dimensional model in all dimensions simultaneously. Also, they can eliminate the
33 need for a three-dimensional device and can minimize or eliminate the need for a user to

1 position operators simultaneously in three dimensions. Hence, a user can think in two
2 dimensions at a time instead of having to think in three dimensions.

3 An interface according to the present invention is illustrated in FIG. 9. The
4 interface can include a processor 100, such as a computer, a two-dimensional input device
5 110, such as a cursor control device, a touchscreen, a trackball, or a mouse, and a display
6 device 120, such as a computer screen. The two-dimensional input device 110 can be
7 operably connected to the processor, as can the display device 120. The display device
8 120 can provide the first view 130 and second view 140 of the object, as discussed above.
9 The interface can also include a memory 150 and a secondary storage 160 for storing data
10 or applications for execution by the processor 100. In particular, processor 100 can
11 execute software routines used to implement those methods shown in FIG. 8.

12 Once a model has been formed, texture mapping of the model can be added to
13 improve the quality of the model. In order to add texture mapping, a two-dimensional
14 texture map can be introduced. Then, the map can be projected onto the model to
15 compute texture coordinates for the voxels. Three-dimensional textures can also be
16 applied according to other methods known in the art.

17 The foregoing detailed description has been given for understanding exemplary
18 implementations of the invention only and no unnecessary limitations should be
19 understood therefrom as modifications will be obvious to those skilled in the art without
20 departing from the scope of the appended claims and their equivalents.